

EUROPEAN PATENT APPLICATION

⑭ Application number: 85106189.5

⑮ Int. Cl.: G 06 F 12/08

⑯ Date of filing: 21.05.85

⑰ Priority: 29.06.84 US 626089

⑱ Date of publication of application: 02.01.86
Bulletin 86/1

⑲ Designated Contracting States: DE FR GB

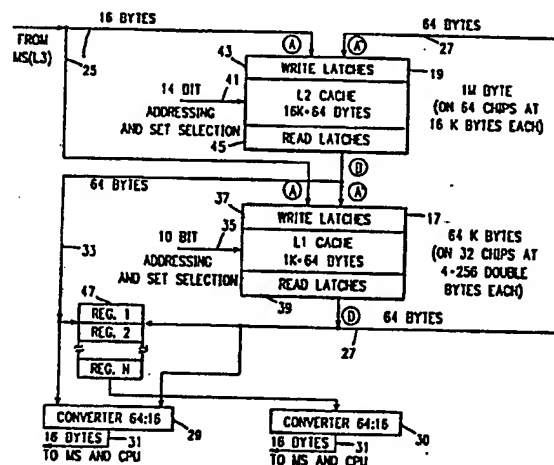
⑴ Applicant: International Business Machines Corporation,
Old Orchard Road, Armonk, N.Y. 10504 (US)

⑵ Inventor: Alchermann, Frederick John, Jr., RD7, Taconic Drive, Hopewell Junction, N.Y. 12533 (US)
Inventor: Blumberg, Rex Harold, 274 Browning Road, Hyde Park New York 12538 (US)
Inventor: Meltzer, David, 68 Maloney Road, Wappingers Falls New York 12590 (US)
Inventor: Pomerene, James Herbert, 403 N. Bedford Road, Chappaqua, N.Y. 10514 (US)
Inventor: Puzak, Thomas Roberts, Route 118 Box 341, RD2, Yorktown Heights, N.Y. 10598 (US)
Inventor: Rechtschaffen, Rudolph Nathan, 24 Innex Road, Scarsdale New York 10583 (US)
Inventor: Sparacio, Frank John, 143-74th Street, No. Bergen, New Jersey 07047 (US)

⑶ Representative: Möhlen, Wolfgang C., Dipl.-Ing., Säumerstrasse 4, CH-8803 Rüschlikon (CH)

⑷ High-speed buffer store arrangement for fast transfer of data.

⑸ In a data processing system comprising multiple cache buffer stores (17, 19) in a hierarchical arrangement, fast transfer of wide data blocks is enabled by particular cache configurations and cache interconnections. On each cache chip, input and output (39, 45) latches are integrated thus avoiding separate intermediate buffering. Input and output latches are interconnected by 64-byte wide data buses (B, A'; D, A'') so that data blocks can be shifted rapidly from one cache hierarchy level to another and back. Chip-internal feedback connections from output to input latches allow to selectively reenter data blocks into a cache after reading. An additional register array (47) is provided so that data blocks after transfer from a cache to main memory or CPU can be subsequently furnished again without accessing the respective cache. The disclosed system allows to transfer wide data blocks within one cycle, thus tying-up caches much less in transfer operations, so that their availability is increased.



HIGH-SPEED BUFFER STORE ARRANGEMENT
FOR FAST TRANSFER OF DATA

This invention is concerned with high-speed buffer stores or caches in data processing systems, and in particular with the design and interconnection of multiple caches to enable fast transfer of data sets between the caches, and also between a cache and the main store or processing unit.

The use of high-speed buffer stores, often called "caches", for improving the operation of data processing systems is well established in the art. Several systems are known in which a plurality of caches are provided.

U.S. Patent 4,141,067 discloses a multiprocessor system in which each CPU has its own cache store. Separate latches are provided between each cache store and its CPU to buffer data. No transfer or interaction between the several caches is provided, as each cache serves its own processor.

In U.S. Patent 4,144,566, a parallel processor is disclosed having a large number of elementary processors connected in parallel. Each elementary processor has its own normal storage unit and its own small capacity fast storage unit. These fast storage units are interconnected to allow the desired parallel processing. However, no transfer of separate data sets between the fast stores or between a selectable fast store and a single common main store are provided.

U.S. Patent 4,228,503 describes a multi-requestor system in which each requestor has its own dedicated cache store. Besides having access to its own cache store for obtaining data, each requestor also has access to all
5 other dedicated cache stores for invalidating a particular data word therein if that same data word has been written by that requestor into its own dedicated cache store. However, a requestor cannot obtain data from another cache which is not its own, and no data transfers
10 between caches are provided.

In U.S. Patent 4,354,232 a computer system is disclosed which has a high-speed cache storage unit. A particular buffer stage is provided between the cache and the main storage and CPU, for storing read and write data
15 transfer commands and associated data. Though flexibility is gained in data transfer, a separate buffer unit and control logic are required solely for this purpose.

The article "Data processing system with second level cache" by F. Sparacio, IBM Technical Disclosure
20 Bulletin, Vol. 21, No. 6, November 1978, pp. 2468-2469, outlines a data processing system having two processors and a two-level cache arrangement between each processor and the common main store. No disclosure is made of the internal organization of the cache stores and of the
25 interconnecting busses and circuits.

An article by S.M. Desai "System cache for high performance processors" which was published in IBM
Technical Disclosure Bulletin, Vol. 23, No. 7A, December 1980, pp. 2915-2917 presents a basic block diagram of a
30 data processing system having plural processors each with

its own dedicated cache store, and a common system cache in a separate level between the dedicated processor caches and main storage. Also in this article, no details are given on interconnecting busses and circuits and on
5 the internal organization of the cache storage units.

It is an object of the invention to devise a high-speed buffer storage arrangement having multiple caches with improved data transfer capabilities between caches and between any cache and the main store or a processor.

10 It is another object to provide a cache buffer organization with improved data transfer capabilities that requires no separate buffer units between the caches or in the data paths.

A further object is to provide a multiple cache
15 buffer system that allows fast transfer of data blocks to and from caches having different access times without the requirement of extra operating cycles for intermediate handling.

The invention for achieving these objects and
20 further advantages is defined in the claims.

The new cache buffer arrangement allows transfer of very large data blocks between storage units within one operation cycle. It is particularly suited for a hierarchical system of high-speed buffers having different
25 speeds and sizes.

Its improved performance is based on special form factors of the internal memory organization, supported by

directly-connected on-chip latches which can be provided with separate external control lines.

Due to the transfer of wide data blocks in parallel mode, the cache stores are tied up in transfer operations much less than it was necessary in systems where several sequential transfers of smaller data blocks are effected. The requirement for wider data paths and associated circuitry is more than compensated by the much higher availability of the cache buffers which is due to the fast, single-operation block transfers.

An embodiment of the invention is described in the sequel with reference to the drawings.

- FIG. 1 is a block diagram of the data flow in a system in which the invention is implemented.
- FIG. 2 shows more details of the two cache stores of FIG. 1 and their interconnections.
- FIG. 3 illustrates the organization of a single chip of the level 1 cache store of FIG. 2, including control and data lines and on-chip latches.
- FIG. 4 illustrates the organization of a single chip of the level 2 cache store of FIG. 2, including control and data lines and on-chip latches.

FIG. 5 shows the addressing structure for selecting a single 64-byte line of data in the level 1 cache store.

FIG. 6 shows the addressing structure for selecting a single 64-byte line of data in the level 2 cache store.

DETAILED DESCRIPTION

(A) STORAGE SYSTEM DATA FLOW

Fig. 1 is a block diagram of the storage system which will be disclosed as an embodiment of the invention. A processor 11 is connected to main storage unit 13 by a storage control unit 15. Two cache high speed buffer stores 17 and 19 are provided to improve the availability of operands and instructions to the processor. The arrangement of the caches in a two-level hierarchy (with the main store being in the highest level L3) brought further improvement, as was e.g. explained in above-mentioned IBM Technical Disclosure Bulletin article by F.J. Sparacio. Cache controls 21 (L1 CTL) and 23 (L2 CTL) are provided for the two cache stores, respectively, and are connected to main storage control unit 15.

Present invention is concerned with the internal organization of the cache buffer stores and their interconnections.

As can be seen from Fig. 1, the level 1 (L1) cache 17 has a capacity of 64 K bytes, and the level 2 (L2)

cache 19 has a capacity of 1M bytes, i.e. L2 is sixteen times as large as L1. Data can be transferred from the main store via 16-byte wide bus 25 to the inputs of both cache buffers. From L1 cache 17, data can be transferred
5 via 64-byte wide bus 27 to a second input of L2 cache 19, and also through a converter 29 to a 16-byte wide bus 31 which is connected to the processor 11 and also through the storage control to main store 13. From L2 cache 19, data can be transferred via 64-byte wide bus 33 to a
10 second input of L1 cache 17, and also through the converter 29 and 16-byte bus 31 to the processor and to the main store.

More details of the two high-speed cache buffers will be disclosed in the following sections.

15 The bus width and storage sizes of this preferred embodiment are of course only one possibility. Other widths and sizes can be selected, depending on the design and application of the respective data processing system.

20 It is also possible to implement the invention in a multiple processor system. In such a multiprocessor system, a single common cache group can be provided between all processors and the common main store, or a separate local group of caches could be devoted to each of the processors with only the main store being commonly used.
25 However, this is immaterial for the invention which is only concerned with the internal organization and interconnection of the multilevel caches, and their interface to the other units of the system.

(B) L1 CACHE, L2 CACHE, AND INTERCONNECTIONS

Fig. 2 shows some more details about the two caches L1 and L2 and their interconnections. Both cache buffers are so organized that data (operands, instructions) can be accessed in portions of 64 bytes, each such portion being designated as a "line" in the following. Thus, one line comprises 64 bytes or 576 bits (each byte including eight data bits and one parity bit, i.e. 1 byte = 9 bits).

Level 1 cache 17 with its capacity of 64 K bytes can hold 1024 (or 1K) lines of 64 bytes each. To select one line location for reading or writing 64 bytes, the cache needs the equivalent of 10 bits which are provided on a group of selection lines 35. Some of these selection bits are used for selecting a set (or subdivision) of the cache, and the others are used for addressing a specific location within the set. This will be explained in more detail in connection with Fig. 3.

L1 cache 17 has write latches 37 which can hold one line or 64 bytes of data. These latches are selectively loaded either from L2 cache via bus 33 (input A') or from main store in four sequential passes via bus 25 (input A). L1 cache 17 further has read latches 39 which also can hold one line = 64 bytes of data. Contents of these latches is furnished to bus 27 (output D).

L1 cache 17 is arranged on 32 integrated circuit chips, each holding four sets of 256 double bytes (as will be shown in more detail in Fig. 3). Of any stored line of 64 bytes, each chip holds one double byte. Thus,

on each of the 32 chips, there are integrated write latches 37 for one double byte (18 bits) and also read latches 39 for one double byte (18 bits).

5 The access time of L1 cache chip is in the order of 3 ns or less.

10 Level 2 cache 19 is of similar but not identical design as L1. With its capacity of 1 M byte it can hold 16,384 (16 K) lines of 64 bytes each. For selecting any one of these lines, the equivalent of 14 selection bits are required which are provided on selection lines 41. Details of selection and addressing in L2 cache 19 will be explained in connection with Fig. 4.

15 L2 cache 19 also has a set of write latches 43 which can hold one line of 64 data bytes. These latches are selectively loaded either from L1 cache via bus 27 (input A") or from main store in four sequential passes via bus 25 (input A) like the L1 cache. L2 cache 19 also has read latches 45 which can hold a line of 64 data bytes. Contents of these latches is furnished to bus 33 (output B).

20 L2 cache 19 is arranged in 64 integrated circuit chips, each holding 16 K single bytes (grouped in sets and subsets, as will be shown in more detail in Fig. 4). Of any stored line of 64 bytes, each chip holds one single byte. Thus, on each of the 64 chips, there are
25 integrated write latches 43 for one byte (9 bits) and also read latches 45 for one byte (9 bits).

 The access time of L2 cache chip 19 is in the order of 20 ns (or less), i.e. much longer as that of L1 cache 17 because of the larger size.

Conv rter 29 receives a 64-byte line from either L1 or L2, and releases it in four successive cycles in 16-byte portions (or sublines) to main store or processor.

Block 47 in Fig. 2 represents an array of N registers which each can hold a 64-byte line which was transferred to converter 29 from either L1 cache or L2 cache. These registers allow to re-use lines of data without accessing again the respective cache high-speed buffer store. The registers feed a second 64:16 converter 30 to allow parallel cache and register readout.

(C) LAYOUT AND CONTROL OF AN L1 CHIP

In Fig. 3, one of the 32 chips constituting the level 1 cache buffer store is shown. This L1 chip 51 comprises four arrays 53, 55, 57, 59 each for storing 256 double bytes (i.e. 256 x 18 bits). It further comprises write latches 37' for storing one double byte (18 bits), and read latches 39' for storing one double byte (18 bits). The 18 bits of write latches 37' are transferred via bus 61 to all four arrays, and bus 63 is provided to transfer 18 bits from any array to read latches 39'. Write and read latches are connected to external buses 25' (input A), 33' (input A'), and 27' (output D), respectively, as was shown in Fig. 2. (Of the total 64-byte capacity of each external bus, only two bytes, i.e. 18 bits are connected to each individual chip 51, as is indicated by the stroke in 25' etc.).

An extra feedback connection 65 is provided on the chip for transferring a double byte from read latches 39' back to write latches 37', thus forming a third input AI to the write latches.

For selecting any one of the 256 double bytes on each array, eight address bits (ADDR L1) are provided on lines 67 and are decoded in decoding circuitry 69. For selecting any one of the four arrays 53, 55, 57, 59, two selection bits (SEL L1) are provided on line 71 and are decoded in decoding circuitry 73 or 74, respectively. The clock signal and the write enabling signal (WRITE L1) on lines 75 are used for array control and timing during a write array operation. In a read operation, four double bytes - one from each of the four arrays - are read simultaneously, and one is gated by selected AND gate circuitry (G) at the end of the array cycle time. The selection is effected by an output signal of decoder 74 which receives the two array selection bits (SEL L1) on lines 71 and which is enabled by a read enabling signal (READ L1) provided on line 77. The signal on line 77 is also used for array control.

Thus, by the ten bits on lines 67 and 71 (which together constitute the selection lines 35 shown in FIG.2), one of the 1024 double bytes stored in the respective chip can be selected. It will be shown in connection with FIG. 5 how these ten addressing/selection bits are developed from a given address.

As there are three inputs to write latches 37', a two-bit control signal "W1" is provided on lines 79 for selecting any one of the inputs A, A' and AI and for enabling write latches 37' to store the two bytes available on the selected input bus.

A further two-bit control signal "WH" is provided on lines 81 to gate either only the left byte or only the right byte of the two bytes available on the selected

input bus, into write latches 37'. This enables selection of individual bytes, or the assembling of two bytes from different sources in a single byte pair.

5 A read control signal "R1" is provided on single-bit line 83 to read latches 39'. This signal when active enables read latches 39' to store the double byte currently available on bus 63, as read from one of the four storage arrays.

10 Control signals W1, WH and R1 (which are furnished by L1 controls 21) are an important feature of the disclosed storage system. They enable to separate internal operation of the chips/cache from external data transfers. Thus, despite different operating speeds or access times of caches L1 and L2 and the main store, direct
15 transfers between the different storage levels are possible with a minimum delay, i.e. without requesting extra storage cycles.

(D) LAYOUT AND CONTROL OF AN L2 CHIP

20 In FIG. 4, one of the 64 chips constituting the level 2 cache buffer store is shown. This L2 chip 91 comprises a large array 93 of 16,384 (16 K) byte positions each holding nine data bits. It further comprises write latches 43' for storing one byte (9 bits) and read latches 45' for storing one byte (9 bits). Bus 95 connects the write latches to array 93, and bus 97 connects
25 the array 93 to the read latches. Write and read latches are connected to external busses 25' (input A), 27' (input A"), and 33' (output B), respectively, as was shown in FIG. 2. (Of the total 64-byte capacity of each external
30 bus, only one byte, i.e. nine bits are connected to

each individual chip 91 as is indicated by the stroke in 25' etc.)

For selecting any one of the 16 K bytes on array 93, twelve address bits (ADDR1 L2, ADDR2 L2) are provided on
5 lines 101 and 103, and two selection bits (SEL L2) on lines 105. (Lines 101, 103 and 105 together constitute the selection lines 41 shown in FIG. 2.) These fourteen bits are decoded in decoding circuitry 107, 109, 111, and the respective signals select a set (or superline) in
10 array 93 and one subset (line) within a selected set. It will be shown in connection with FIG. 6 how the addressing/selection bits are developed from a given address.

Additional lines 113 and 115 are provided for furnishing a write enabling signal (WRITE L2) and a read enabling signal (READ L2), respectively, to storage array
15 93.

A two-bit control signal "W2" is provided to write latches 43' on lines 117 for selecting one of the two inputs A and A" and for enabling write latches 43' to
20 store the single byte available on the selected input bus.

A read control signal "R2" is provided to read latches 45 on single-bit line 119.

This signal when active enables read latches 45' to
25 store the single byte currently available on bus 97 as read from storage array 93.

Control signals W2 and R2 (which are furnished by L2 controls 23) are an important feature of the disclosed

storage system, in connection with the on-chip write and read latches, because these features significantly enhance the inter-level transfer capabilities of the cache storage hierarchy (as was already mentioned at the
5 end of the previous section).

(E) ADDRESSING OF L1 CACHE

FIG. 5 illustrates how the addressing/selection signals for level 1 cache buffer store 17 are developed from a given address. The 27 bits of a virtual address
10 are stored in register 121. The lowest-order 6 bits are used for selecting one byte of a 64-byte line read from the L1 cache. All other bits are used for addressing one 64-byte line in cache.

A directory look-aside table (DLAT) 123 is provided
15 for storing recently translated addresses, as is well-known in virtual storage systems. The DLAT is subdivided into 256 congruence classes. All virtual addresses in which bits 7...14 are identical form one congruence class, or associative set. Thus, these eight bits are
20 used to select the respective congruence class (or row) in the DLAT. Each congruence class has two entries 125, each of them storing a "STO" address field (17 bits), a virtual address field (7 bits) and the corresponding translated absolute address field (15 bits). Now when a
25 congruence class was selected, the seventeen bits of a given "STO" address and the seven highest-order bits 0...6 of the virtual address register are compared with the respective fields in the two DLAT entries. If no match occurs, a translation must be made and entered into
30 DLAT. If a match occurs, the respective translated

fifteen absolute address bits are furnished at the DLAT output.

For addressing the cache and its directory, also congruence classes are used which are different from the
5 DLAT congruence classes. For the cache, all virtual addresses in which bits 13...20 are identical form one congruence class or associative set. These eight bits are transferred to L1 directory 127 and L1 cache 17 for selecting one congruence class (or row) of 256. The
10 directory as well as the cache are 4-set associative, i.e. they have four entries per congruence class or row. In the directory, each entry 129 holds a 15-bit absolute address; in the cache, each entry 131 holds a whole data line of 64 bytes.

15 The fiveteen address bits furnished by the DLAT are compared in the L1 directory with all four entries of the selected row. If no match occurs (cache miss), the respective line must be fetched to cache and the address entered into the directory. If a match occurs (cache
20 hit), a two-bit signal identifying the respective set (column) is transferred to the L1 cache for selecting there the corresponding set (column).

Now the eight addressing bits and the two set selection bits are available on lines 67 and 71 of the
25 cache, respectively, and can be used for selecting a double byte on each of the 32 cache chips, as was explained in connection with FIG. 3. The 64-byte line is then stored in the read latches of all chips, and becomes available on output bus 27.

(F) ADDRESSING OF L2 CACHE

FIG. 6 shows how the addressing/selection signals for level 2 cache buffer store 19 are developed from a given address. It is assumed that the virtual address was already translated into a 27-bit absolute address which is stored in a register 133. The twelve low-order bits 15...26 are taken directly from the virtual address whereas the 15 high-order bits 0...14 are obtained from a directory look-aside table DLAT, as was explained for L1 cache in connection with FIG. 5.

The six lowest-order bits 21...26 of the absolute address are used for selecting one byte of a 64-byte line read from the L2 cache. All other bits (0...20) are used for addressing one 64-byte line in cache.

The level 2 cache and its directory are also subdivided into congruence classes. The nine bits 7...15 of the absolute address determine the congruence class so that 512 classes can be distinguished.

L2 directory 135 has 512 rows (for the 512 congruence classes) each comprising four entries 137 (4-way associativity). Thus $4 \times 512 = 2,048$ data sets can have their address in the L2 directory. Each such data set is a superline comprising eight 64-byte lines stored in 64 chips in cache.

Addressing of a superline is as follows: The nine bits (7...15) determining the congruence class select one row in the L2 directory. Nine further bits of the absolute address (bits 0...6 and 16 and 17) which identify the superline (8 lines) are furnished to the directory and

are compared with the four 9-bit entries in the selected row. If no match occurs (cache miss), a fetch in main store must be made and the directory updated. If a match occurs (cache hit), then the respective column is identified by a bit pair furnished at the output of L2 directory 135. This bit pair determines where within the respective congruence class the addressed superline is located in cache.

L2 cache 19 receives the nine bits determining the congruence class (which could be designated as "row" in cache) on lines 101, and it receives the four bits determining the set or superline within that congruence class (or row) on lines 105.

To finally select a single 64-byte line 139 within the superline, three absolute address bits (18...20) are furnished to L2 cache on lines 103. Thus, fourteen bits are available at the inputs of the cache to select one 64-byte line out of the totally stored 16 K lines. Each of the 64 chips of the L2 cache furnishes one byte (9 bits) of the selected line, and all 64 bytes appear simultaneously on output bus 33.

For writing into the caches, the same addressing mechanism is used as described above for reading.

C L A I M S

1. Apparatus for the efficient transfer of a multiple byte data entity between caches in a storage system comprising a plurality of cache storage units (17, 19), and between said cache storage (13) units and main storage and a processor (11) in a data processing system, characterized in that

- each of said cache storage units has integrated therewith a set of input latches (37, 43) and a set of output latches (39, 45), and that the outputs (B, D) of said output latches are connected to the inputs (A', A") of said input latches of at least one other of said cache storage units.

2. Apparatus in accordance with claim 1, characterized in that said cache storage units (17, 19) are arranged in a multilevel storage hierarchy, that said output latches of each cache storage unit are connected to the input latches of the cache storage unit on the next level, with the output latches of the lowest-level cache being connected to the input latches of the highest-level cache.

3. Apparatus in accordance with claim 1 or 2, characterized in that the input latches (37, 43) of all cache storage units (17, 19) are connected to a common main storage data bus (25), and that the output latches (39, 45) of all cache storage units are connected to a selective gating means for selective transfer of a multiple byte data entity to the processor and/or to the main storage data bus.

4. Apparatus in accordance with claim 1 or 3, characterized in that said output latches of said cache storage units are also connected to a register array (47) for selectively storing therein a multiple byte data entity that is transferred to said processor and/or to said main store.

5. In a data processing system, high-speed buffer storage apparatus connected to a main store (13) and to a processor (11), said buffer storage apparatus comprising a plurality of cache buffer stores (17, 19) each consisting of a plurality of storage chips (51, 91), characterized in that

- in each said cache buffer store, each chip includes a set of write latches (37, 43) and a set of read latches (39, 45) for holding data to be written into or read from storage circuits on said chip, respectively,

- bus interconnections (B, A'; D, A'') are provided between the read latches of all chips of at least one cache buffer store and the write latches of all chips of at least one other cache buffer store, for parallel transfer of a multi-byte data block between said cache buffer stores.


6. Apparatus in accordance with claim 5, characterized in that for each of said chips (51; 91), separate write and read control lines (79, 81, 83; 117, 119) are provided for said write and read latches (37, 39; 43, 45), respectively, so that data can be loaded into said write or read latches independently of a write or read operation in the storage circuits (53, 55, 57, 59; 93) on the respective chip.

7. Apparatus in accordance with claim 5, characterized in that at least two cache stores (17, 19) are provided constituting a two-level cache buffer storage hierarchy, and that the outputs of the read latches of
5 each cache store are connected to the write latches of the other cache store.

8. Apparatus in accordance with claim 5 or 7, characterized in that the interconnections (B, A'; D, A") between said cache storage units have a width for transferring w data bytes in parallel; that the inputs of the
10 write latches (37, 43) of all cache stores are additionally connected to a common main storage data bus (25) having a width which is only a fraction w/n of the inter-cache bus width, and that the outputs of the read
15 latches (39, 45) of all cache stores are connected by at least one output data bus (27) having a width of w data bytes, to converter means (29) for sequentially furnishing fractions of a w-byte data block to a processor and/or main store data bus (31) having only the fractional width w/n .
20

9. Apparatus in accordance with claim 8, characterized in that an additional register array (47) having a capacity of at least w data bytes is connected to said output data bus (27), for storing at least one data block
25 of w bytes which was furnished to said converter means (29), and that second converter means (30) is provided for sequentially furnishing fractions of any w-byte data block stored in said register array to a processor and/or main store data bus (31).

10. Apparatus in accordance with claim 5, characterized in that in at least one of said cache stores a feedback connection (65) is provided on each chip (51) between the outputs of the read latches (39) and the
5 inputs of the write latches (37) of the respective chip.




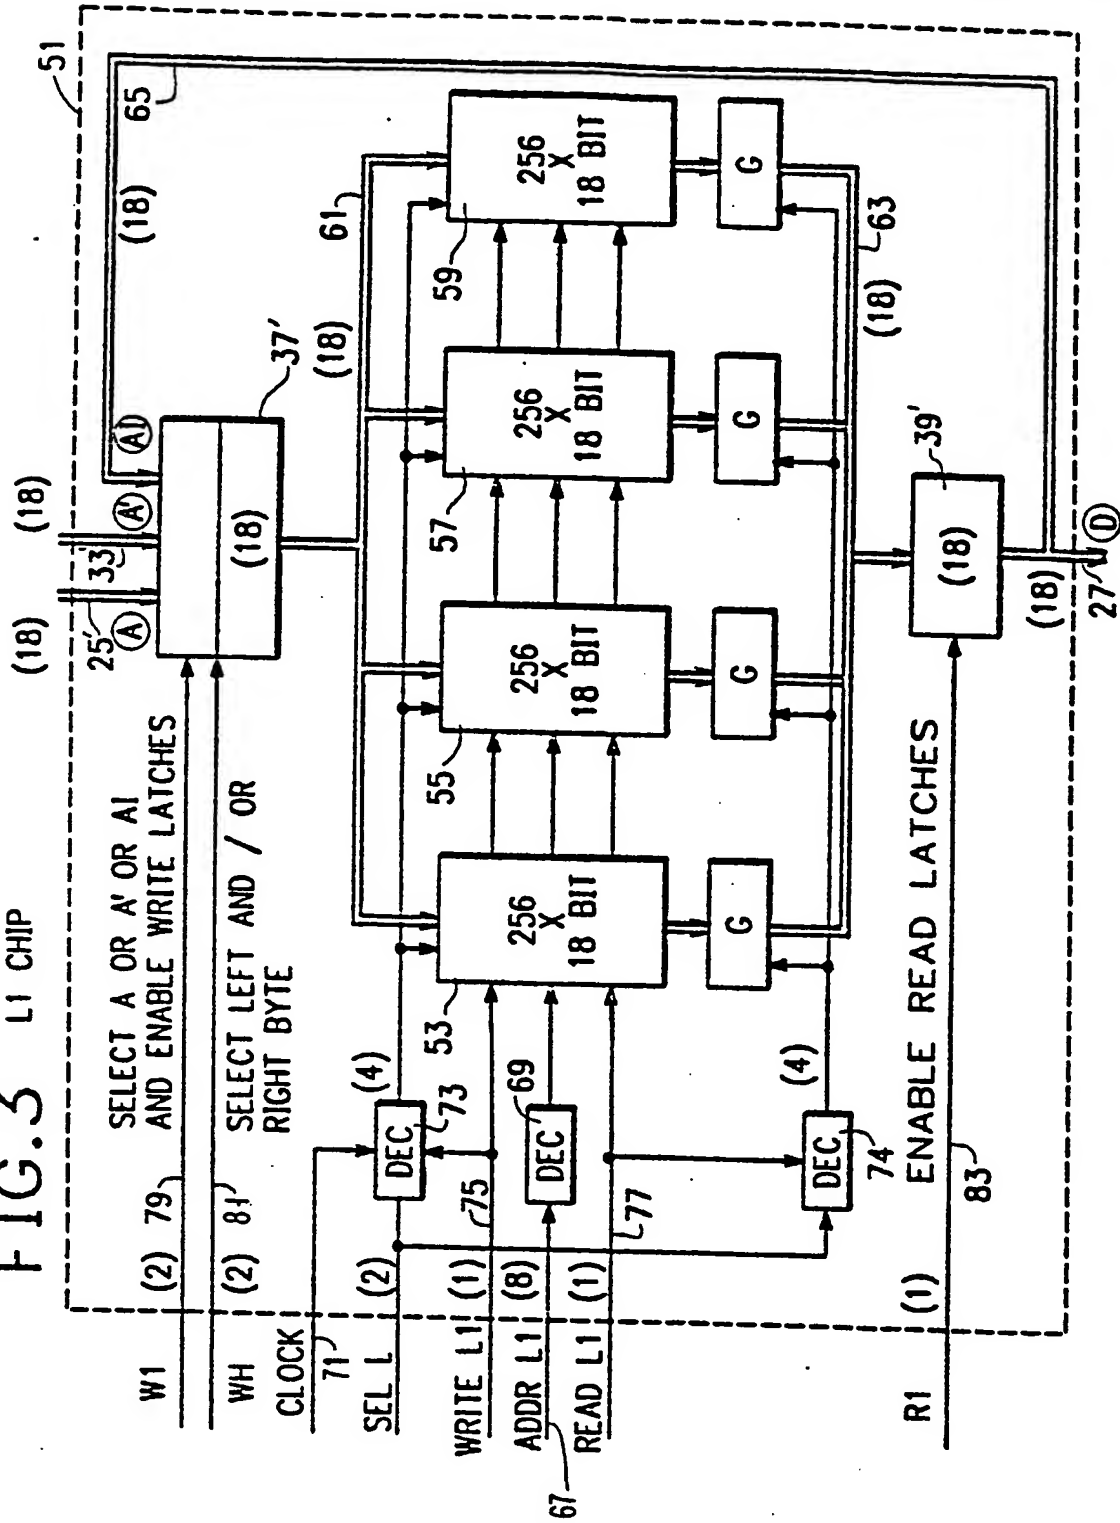




FIG. 2


FIG. 3 L1 CHIP






0166192
FIG. 5







0 166 192
A3

EUROPEAN PATENT APPLICATION

⑤ Int. Cl.⁴: G 06 F 12/08

② Date of filing: 21.05.85

⑦ Applicant: International Business Machines Corporation, Old Orchard Road, Armonk, N.Y. 10504 (US)

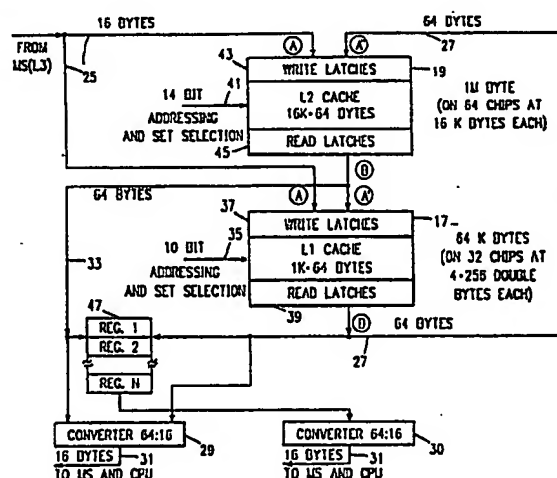
(72) Inventor: Aichelmann, Frederick John, Jr., RD7, Taconic Drive, Hopewell Junction, N.Y. 12533 (US)
Inventor: Blumberg, Rex Harold, 274 Browning Road, Hyde Park New York 12538 (US)
Inventor: Meltzer, David, 68 Maloney Road, Wappingers Falls New York 12590 (US)
Inventor: Pomerene, James Herbert, 403 N.Bedford Road, Chappaqua, N.Y. 10514 (US)
Inventor: Puzak, Thomas Roberts, Route 118 Box 341, RD2, Yorktown Heights, N.Y. 10598 (US)
Inventor: Rechtschaffen, Rudolph Nathan, 24 Innex Road, Scarsdale New York 10583 (US)
Inventor: Sparacio, Frank John, 143-74th Street, No. Bergen, New Jersey 07047 (US)

⑦ Representative: Möhlen, Wolfgang C., Dipl.-Ing., IBM Corporation Säumerstrasse 4, CH-8803 Rüschlikon (CH)

⑧ Date of deferred publication of search report: 26.04.89 Bulletin 89/17

⑤4 · High-speed buffer store arrangement for fast transfer of data.

57) In a data processing system comprising multiple cache buffer stores (17, 19) in a hierarchical arrangement, fast transfer of wide data blocks is enabled by particular cache configurations and cache interconnections. On each cache chip, input and output (39, 45) latches are integrated thus avoiding separate intermediate buffering. Input and output latches are interconnected by 64-byte wide data buses (B, A'; D, A'') so that data blocks can be shifted rapidly from one cache hierarchy level to another and back. Chip-internal feedback connections from output to input latches allow to selectively reenter data blocks into a cache after reading. An additional register array (47) is provided so that data blocks after transfer from a cache to main memory or CPU can be subsequently furnished again without accessing the respective cache. The disclosed system allows to transfer wide data blocks within one cycle, thus tying-up caches much less in transfer operations, so that their availability is increased.





European Patent
Office

EUROPEAN SEARCH REPORT

0166192

Application Number

EP 85 10 6189

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int. Cl. 4)
X	EP-A-0 100 943 (IBM) * Page 2, line 30 - page 5, line 30; figure; page 7, lines 7-15 *	1,3-10	G 06 F 12/08
Y	---	2	
Y	GB-A-2 107 092 (WESTERN ELECTRIC) * Figure 1; page 3, lines 10-61 *	2	
A	---	1,3-10	
A	ELECTRONIC ENGINEERING, vol. 53, no. 650, Mid-March 1981, pages 27-30, London, GB; J. SCARISBRICK: "Large scale multi-port memories permit asynchronous operation" * Figure 1; page 28, lines 15-26; figure 3; page 30, lines 20-32 * -----	1,5	
			TECHNICAL FIELDS SEARCHED (Int. Cl. 4)
			G 06 F 12/08 G 11 C 8/00 G 06 F 5/06
The present search report has been drawn up for all claims			
Place of search THE HAGUE		Date of completion of the search 06-02-1989	Examiner MASCHE C.M.
CATEGORY OF CITED DOCUMENTS			
X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document		I : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons ----- & : member of the same patent family, corresponding document	